Volume Title: A Theory of the Consumption Function

Volume Author/Editor: Milton Friedman

Volume Publisher: Princeton University Press

Volume ISBN: 0-691-04182-2

Volume URL: http://www.nber.org/books/frie57-1

Publication Date: 1957

Chapter Title: The Permanent Income Hypothesis

Chapter Author: Milton Friedman

Chapter URL: http://www.nber.org/chapters/c4405

Chapter pages in book: (p. 20 - 37)

CHAPTER III

The Permanent Income Hypothesis

THE magnitudes termed "permanent income" and "permanent consumption" that play such a critical role in the theoretical analysis cannot be observed directly for any individual consumer unit. The most that can be observed are actual receipts and expenditures during some finite period, supplemented, perhaps, by some verbal statements about expectations for the future. The theoretical constructs are *ex ante* magnitudes; the empirical data are *ex post*. Yet in order to use the theoretical analysis to interpret empirical data, a correspondence must be established between the theoretical constructs and the observed magnitudes.

The most direct way to do so, and the one that has generally been followed in similar contexts, is to construct estimates of permanent income and permanent consumption for each consumer unit separately by adjusting the cruder receipts and expenditure data for some of their more obvious defects, and then to treat the adjusted *ex post* magnitudes as if they were also the desired *ex ante* magnitudes. Cash expenditures during a particular time period that are regarded as expenses of earning income can be deducted from cash receipts during the corresponding time period; accrual methods of accounting can be substituted for cash accounting for some or all income items; expenditures on durable consumer goods can be regarded as capital expenditures and only the imputed value of services rendered included as consumption; and so on. These adjustments clearly reduce the difference between the statistical estimates and the theoretical constructs and are therefore highly desirable. But even when they are carried as far as is at all feasible, the resulting magnitudes, interpreted as estimates of permanent income and permanent consumption, are not consistent with equation (2.6): measured consumption turns out to be a smaller fraction of measured income for high than for low measured incomes even for groups of consumer units for whom it does not seem reasonable to attribute this result to differences in the values of i, w, or u.

We are thus driven either to reject equation (2.6), which is what earlier workers have done, or to resort to more indirect means of establishing a correspondence between the theoretical constructs and the observed magnitudes, which is what I propose to do. One indirect means is to use evidence for other time periods and other consumer units to interpret data for one consumer unit for one period. For example, if Mr. A's measured income fluctuates widely from year to year while Mr. B's is highly stable, it seems reasonable that Mr. A's measured income is a poorer index of his permanent income than Mr. B's is of his. Again, suppose Mr. A's measured income in any period is decidedly lower than the average measured income of a group of individuals who are similar to him in characteristics that we have reason to believe affect potential earnings significantly—for example, age, occupation, race, and location. It then seems reasonable to suppose that Mr. A's measured income understates his permanent income.

The following formalization of the relation between the theoretical constructs and observed magnitudes is designed to facilitate the use of such evidence. Its central idea is to interpret empirical data as observable manifestations of theoretical constructs that are themselves regarded as not directly observable.

1. *The Interpretation of Data on the Income and Consumption of Consumer Units*

Let y represent a consumer unit's measured income for some time period, say a year. I propose to treat this income as the sum of two components: a permanent component (y_p), corresponding to the permanent income of the theoretical analysis, and a transitory component (y_t),[1] or

$$y = y_p + y_t . \tag{3.1}$$

The permanent component is to be interpreted as reflecting the effect of those factors that the unit regards as determining its capital value or wealth: the nonhuman wealth it owns; the personal attributes of the earners in the unit, such as their training, ability, personality; the attributes of the economic activity of the earners, such as the occupation followed, the location of the economic activity, and so on. It is analogous to the "expected" value of a probability distribution. The transitory component is to be interpreted as reflecting all "other" factors, factors that are likely to be treated by the unit affected as

[1] The terminology, and much of the subsequent analysis, is taken from Friedman and Kuznets, *Income from Independent Professional Practice*, pp. 325–38, 352–364.

"accidental" or "chance" occurrences, though they may, from another point of view, be the predictable effect of specifiable forces, for example, cyclical fluctuations in economic activity.[2] In statistical data, the transitory component includes also chance errors of measurement; unfortunately, there is in general no way to separate these from the transitory component as viewed by the consumer unit.

Some of the factors that give rise to transitory components of income are specific to particular consumer units, for example, illness, a bad guess about when to buy or sell, and the like; and, similarly, chance errors of measurement. For any considerable group of consumer units, the resulting transitory components tend to average out, so that if they alone accounted for the discrepancies between permanent and measured income, the mean measured income of the group would equal the mean permanent component, and the mean transitory component would be zero. But not all factors giving rise to transitory components need be of this kind. Some may be largely common to the members of the group, for example, unusually good or bad weather, if the group consists of farmers in the same locality; or a sudden shift in the demand for some product, if the group consists of consumer units whose earners are employed in producing this product. If such factors are favorable for any period, the mean transitory component is positive; if they are unfavorable, it is negative.[3] Similarly, a systematic bias in measurement may produce a nonzero mean transitory component in recorded data even though the transitory factors affecting consumer units have a zero effect on the average.

Similarly, let c represent a consumer unit's expenditures for some time period, and let it be regarded as the sum of a permanent component (c_p) and a transitory component (c_t), so that

$$c = c_p + c_t \, . \tag{3.2}$$

Again, some of the factors producing transitory components of consumption are specific to particular consumer units, such as unusual sickness, a specially favorable opportunity to purchase, and the like; others affect groups of consumer units in the same way, such as an

[2] This division is, of course, in part arbitrary, and just where to draw the line may well depend on the particular application. Similarly, the dichotomy between permanent and transitory components is a highly special case. See *ibid.*, pp. 352–364, for a generalization to a larger number of components.

[3] Note the difference from *ibid.*, p. 326, where the mean transitory component can be taken to be zero without loss of generality. The difference reflects a narrower definition of transitory component in *ibid.* plus the use of the concept to compare the same group in two years.

unusually cold spell, a bountiful harvest, and the like. The effects of the former tend to average out; the effects of the latter produce positive or negative mean transitory components for groups of consumer units; the same is true with chance and systematic errors of measurement.

It is tempting to interpret the permanent components as corresponding to average lifetime values and the transitory components as the difference between such lifetime averages and the measured values in a specific time period. It would, however, be a serious mistake to accept such an interpretation, for two reasons. In the first place, the experience of one unit is itself but a small sample from a more extensive hypothetical universe, so there is no reason to suppose that transitory components average out to zero over the unit's lifetime. In the second place, and more important, it seems neither necessary nor desirable to decide in advance the precise meaning to be attached to "permanent." The distinction between permanent and transitory is intended to interpret actual behavior. We are going to treat consumer units *as if* they regarded their income and their consumption as the sum of two such components, and *as if* the relation between the permanent components is the one suggested by our theoretical analysis. The precise line to be drawn between permanent and transitory components is best left to be determined by the data themselves, to be whatever seems to correspond to consumer behavior.

Figure 2 is designed to bring out more explicitly the wide range of possible interpretations of permanent income. This figure refers to a single consumer unit, the head of which is assumed to be 30 years of age on the date in 1956 for which the figure is drawn. We may suppose the unit to have been formed when the head was aged 20. Measured income experience from 20 to 30, as recorded in the solid jagged line, is a datum; so also, of course, are other items not recorded in the figure, such as the amount of nonhuman wealth possessed, the occupation of the head and of other members of the unit, location, and so on. Future measured income experience is uncertain. The scatter of dots for later ages is intended to represent the possibilities as *viewed* by the unit; for each future date, there is some anticipated probability distribution of measured income. Because of the limitations of a two-dimensional figure, this scatter diagram seriously misrepresents the situation in one important respect. It suggests that the probability distributions at different ages are independent, whereas in general they might be expected to be interdependent. The distribution anticipated for age 40, for example, if a high measured income is realized at age 31 would presumably be different from the distribution expected if a low measured income is realized.[4] But this defect of the figure is not serious for our present limited purpose.

The scatter in the figure should not be confused conceptually with the corresponding scatter that would be generated by plotting the contemporaneous incomes of a large number of units with heads of different age. The scatter in the figure is the anticipated experience of one unit, not the realized experience of many. In forming its

FIGURE 2

Illustration of Alternative Interpretations of Permanent Income

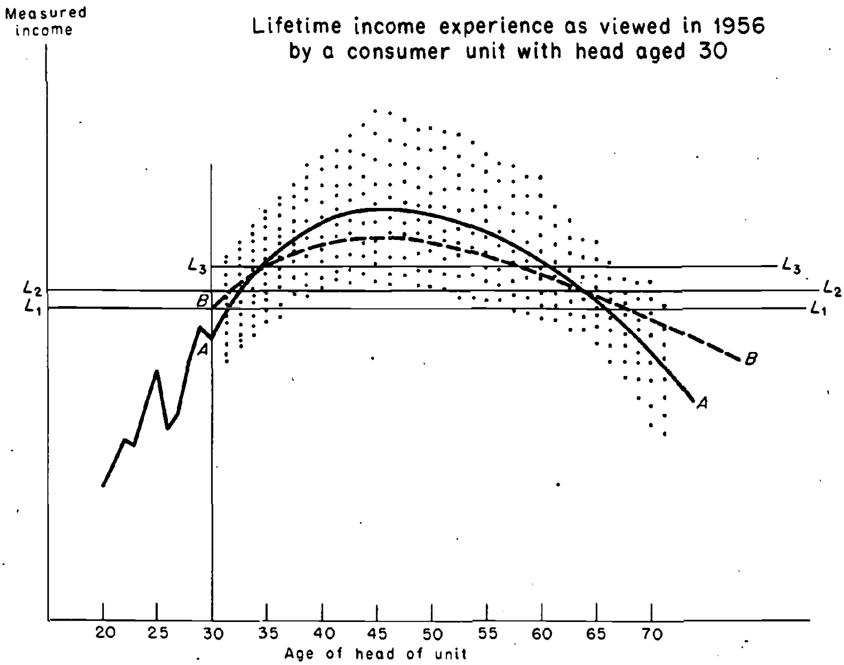

anticipations, the one unit may well take into account the contemporaneous experience of units which are of different age but alike in respect of other factors such as occupation, nonhuman wealth, etc.; and it may for some purposes usefully be regarded as doing so by simply accepting the contemporaneous differences as describing its own future possibilities. On the conceptual level, however, there is no

[4] The most general description would be in terms of a probability distribution of alternative age-measured-income functions. It should be noted that the generalized analytical formulation in *ibid.*, pp. 352–364, allows fully for interdependence.

need to foreclose the possibility that the unit will take other information into account as well.[5]

The solid curve AA in Figure 2 is the mean of the probability distributions anticipated for future years. It is one possible interpretation of the permanent income to which consumption is adapted. The horizontal line L_1L_1 is the mean lifetime income as anticipated at age 20; L_2L_2, as anticipated at age 30, taking into account realized experience from age 20 to age 30; L_3L_3, mean income anticipated at age 30 for the remaining lifetime of the unit. Each of these is another possible interpretation of permanent income, and almost at the opposite extreme of the spectrum from AA. No one of these has very great intuitive appeal as *the* permanent income to which consumption is adapted; AA, because it implies an exceedingly short time horizon; L_1L_1, L_2L_2, and L_3L_3, not only because they imply an extremely long time horizon, but also because they imply that units can borrow on the basis of anticipated receipts from both human and nonhuman wealth at the same interest rate at which they can lend accumulated nonhuman wealth. L_3L_3 has the further objection that it supposes no carry-over into the present of past adaptations. The dashed curve BB is an intermediate interpretation, intended to be something of an average of AA and L_1L_1 or L_2L_2. Something like this seems intuitively the most plausible interpretation, but intuitive plausibility gives little guidance to the exact kind of average, or length of horizon. For this, we must rely on the empirical evidence. (See Chapter VII, where a tentative estimate is made on the basis of existing evidence.)

Figure 2 is drawn for a particular date. There is nothing about the concept of permanent income that requires the relevant parts of the figure to remain the same for any later date. Aside from the point already made, that in advance the probability distribution for any future date depends on the measured income actually experienced, the whole joint probability distribution may be shifted by occurrences that were entirely unanticipated at the date in question. In our empirical work, we shall sometimes find it desirable to suppose that permanent income, or the age pattern of permanent income, remains unchanged over a period of years, but it should be clear that this is an empirical specialization of a more general concept.

2. *A Formal Statement of the Permanent Income Hypothesis*

In its most general form our hypothesis about the consumption function, which we shall hereafter refer to as the permanent income

[5] For example, in sec. 3 of Chap. IV below, the unit is interpreted as modifying contemporaneous experience by information on the secular trend of income.

hypothesis, is given by the three equations (2.6), (3.1), and (3.2):

$$c_p = k(i, w, u)y_p, \tag{2.6}$$

$$y = y_p + y_t, \tag{3.1}$$

$$c = c_p + c_t. \tag{3.2}$$

Equation (2.6) defines a relation between permanent income and permanent consumption. It specifies that the ratio between them is independent of the size of permanent income but does depend on other variables, in particular: (1) the rate of interest (i) or sets of rates of interest at which the consumer unit can borrow or lend; (2) the relative importance of property and nonproperty income, symbolized by the ratio of nonhuman wealth to income (w); and (3) the factors symbolized by the portmanteau variable u determining the consumer unit's tastes and preferences for consumption versus additions to wealth. The most significant of the latter factors probably are (a) the number of members of the consumer unit and their characteristics, particularly their ages, and (b) the importance of transitory factors affecting income and consumption, measured, for example, by the "spread" or standard deviation of the probability distributions of the transitory components relative to the size of the corresponding permanent components. Equations (3.1) and (3.2) define the connection between the permanent components and the measured magnitudes.

In this most general form the hypothesis is empty, in the sense that no empirical data could contradict it. Equations (3.1) and (3.2) are purely definitional; they add two equations but also two additional unknowns, the transitory components. There are a variety of ways to specialize the hypothesis so that it is capable of being contradicted by observed data. The one I shall use is to specify some of the characteristics of the probability distributions of the transitory components. A particularly simple specification, yet one that seems adequate to explain existing evidence, is to suppose that the transitory components of income and consumption are uncorrelated with one another and with the corresponding permanent components, or

$$\rho_{y_t y_p} = \rho_{c_t c_p} = \rho_{y_t c_t} = 0, \tag{3.3}$$

where ρ stands for the correlation coefficient between the variables designated by the subscripts.

The assumptions that the first two correlations in (3.3)—between the permanent and transitory components of income and of consumption—are zero seem very mild and highly plausible. Indeed, by themselves, they have little substantive content and can almost be regarded as simply completing or translating the definitions of transitory and permanent components; the qualitative notion that the transitory component is intended to embody is of an accidental and transient addition to or subtraction from income, which is almost equivalent to saying an addition or subtraction that is not correlated with the rest of income. The merging of errors of measurement with transitory components contributes further to the plausibility that these correlations are zero.

For a group of individuals, it is plausible to suppose that the absolute size of the transitory component varies with the size of the permanent component: that a given random event produces the same percentage rather than the same absolute increase or decrease in the incomes of units with different permanent components. This may make more convenient an alternative definition of transitory component that is suggested below; it is not, however, inconsistent with zero correlation. Zero correlation implies only that the *average* transitory component—the algebraic average in which positive and negative components offset one another—is the same for all values of the permanent component. For example, suppose that the transitory component is equally likely to be plus or minus 10 per cent of the permanent component. The average transitory component is then zero for all values of the permanent component, although the average absolute value, which disregards the sign of the components, is directly proportional to the permanent component.

The plausibility of taking our definition of transitory components to imply a zero correlation for a group of consumer units depends somewhat on the criteria determining membership in the group. The clearest example is a classification of units by the size of their measured income. For each such group, the correlation between permanent and transitory components is necessarily negative, since with a common measured income the permanent component can be relatively high only if the transitory component is relatively low, and conversely.[6]

The assumption that the third correlation in (3.3)—between the transitory components of income and consumption—is zero is a much stronger assumption. It is primarily this assumption that introduces important substantive content into the hypothesis and makes it susceptible of contradiction by a wide range of phenomena capable of being observed. The ultimate test of its acceptability is of course whether such phenomena are in fact observed, and most of what follows is devoted to this question. It is hardly worth proceeding to such more refined tests, however, unless the assumption can pass—or at least not fail miserably—the much cruder test of consistency

[6] See *ibid.*, pp. 326 and 327.

with casual observation of one's self and one's neighbors, so some comments on the intuitive plausibility of the assumption are not out of order.

The common notion that savings, or at least certain components of savings, are a "residual" speaks strongly for the plausibility of the assumption. For this notion implies that consumption is determined by rather long-term considerations, so that any transitory changes in income lead primarily to additions to assets or to the use of previously accumulated balances rather than to corresponding changes in consumption.

Yet from another point of view, the assumption seems highly implausible. Will not a man who receives an unexpected windfall use at least some part of it in "riotous living," i.e. in consumption expenditures? Would he be likely to add the whole of it to his wealth? The answer to these questions depends greatly on how "consumption" is defined. The offhand affirmative answer reflects in large measure, I believe, an implicit definition of consumption in terms of purchases, including durable goods, rather than in terms of the value of services. If the latter definition is adopted, as seems highly desirable in applying the hypothesis to empirical data—though unfortunately I have been able to do so to only a limited extent—much that one classifies offhand as consumption is reclassified as savings. Is not the windfall likely to be used for the purchase of durable goods? Or, to put it differently, is not the timing of the replacement of durable goods and of additions to the stock of such goods likely to some extent to be adjusted so as to coincide with windfalls?

Two other considerations argue for the plausibility of the assumption that transitory components of income and consumption are uncorrelated. First, the above identification of a windfall with transitory income is not precise. Suppose, for example, inheritances are included in a particular concept of measured income. Consider a consumer unit whose receipts remain unchanged over a succession of time periods except that it receives an inheritance in the final period. If the inheritance was expected to occur some time or other, it will already have been allowed for in permanent income; the transitory component of income is only the excess of the inheritance over this element of permanent income. There seems no reason why the receipt of the inheritance should make consumption in the final period different from that of preceding periods, except through inability to borrow in advance on the strength of the inheritance. But this implies that the receipt of the inheritance changes w (the ratio of wealth to income) in (2.6); it is therefore already taken into account in the hypothesis. There is no essential difference if the inheritance is unexpected. The effect of the inheritance is then to increase the permanent income of the unit, and this will justify a higher consumption in the final period; again the transitory component is only the excess of the windfall over this element of permanent income, and it is no longer intuitively obvious that it should lead to an increase in current consumption.[7] The second consideration is that just as there are instances in which one would expect a transitory increase in income to produce a transitory increase in consumption, so also there are instances in which one would expect the reverse. The simplest example is when a transitory increase in income reduces opportunities for consumption as when it is obtained by working longer hours or going to a backward country. Such negative and positive correlations will tend to offset one another.

The preceding remarks abstract from errors of measurement. Yet, as noted, in any statistical analysis errors of measurement will in general be indissolubly merged with the correctly measured transitory components. The effect on the correlation between statistically recorded transitory components of income and consumption depends critically on how the statistical data are obtained. If income and consumption are measured independently, the errors of estimate might be expected to be independent as well and therefore to contribute toward a small or zero observed correlation between transitory components of income and consumption. On the other hand, if consumption is estimated, as it frequently is, by measuring independently savings and income and subtracting the former from the latter, then measured consumption and measured income have common errors of measurement. This tends toward a positive observed correlation between transitory components of income and consumption.

The purpose of these remarks is not to demonstrate that a zero correlation is the *only* plausible assumption—neither evidence like that alluded to nor any other can justify such a conclusion. Its purpose is rather to show that common observation does not render it absurd to suppose that a hypothesis embodying a zero correlation can yield a fairly close approximation to observed consumer behavior. The assumption that the correlation between transitory components of income and consumption is zero could, of course, be replaced by the less restrictive assumption that it is a positive number between zero and unity, but this would greatly weaken the hypothesis and reduce its potential usefulness for predicting behavior. It seems highly undesirable to do so until and unless a significant contradiction

[7] I owe this point to Modigliani and Brumberg, "Utility Analysis and the Consumption Function," pp. 405–406.

arises between the stronger hypothesis and empirical evidence on consumer behavior.

A particularly simple special case of the hypothesis arises if, in addition to (3.3), it is assumed that the mean transitory components of consumption and income are zero, or

(3.4) $$\mu_{y_t} = \mu_{c_t} = 0\,,$$

where μ stands for the mean of the variable designated by its subscript. This assumption is eminently reasonable if the probability distribution in question is sufficiently comprehensive. In general, however, we shall want to use conditional probability distributions, for example, the distribution of transitory components in a particular year, or for members of a particular group. In such cases, it will generally be undesirable to assume that (3.4) holds, just as for the single consumer unit viewed *ex post* it is undesirable to assume that the transitory components themselves are necessarily zero.

It may be desirable or necessary to impose additional conditions on the probability distributions to facilitate the estimation of the parameters of the system from observed data. I shall, however, largely neglect the problem of statistical estimation, and so we need not go into such conditions.

A more important qualification is that, for simplicity of exposition, equations (3.1) and (3.2) express the relation between observed income and its permanent and transitory components as additive. The form of the relation is important because it may affect the empirical validity of such specifications of the characteristics of the probability distributions as (3.3) and (3.4), as well as the validity of using specifications of other characteristics of the distribution that are convenient statistically. From this point of view, I conjecture that a multiplicative specification is preferable for income and consumption data. If we let capital letters stand for the logarithms of the variables designated by the corresponding lower case letters, the equations defining the hypothesis then take the following alternative form:

(2.6′) $$C_p = K(i, w, u) + Y_p,$$

(3.1′) $$Y = Y_p + Y_t,$$

(3.2′) $$C = C_p + C_t,$$

(3.3′) $$\rho_{Y_t Y_p} = \rho_{C_t C_p} = \rho_{Y_t C_t} = 0.$$

Many of the results that follow apply equally to both forms of the hypothesis, requiring only that the same symbol be interpreted in one case as an absolute value, in the other, as a logarithm. For any significant results for which this is not true, the logarithmic expressions are given in footnotes.

3. *The Relation between Measured Consumption and Measured Income*

Suppose we have observations on consumption and income for a number of consumer units, for all of whom the k of equation (2.6) can be taken to be numerically the same. Let us proceed, as is usually done in family budget studies, to estimate from these data a relation between consumption and income. For simplicity, let the relation to be estimated be linear, say:

$$c = \alpha + \beta y, \tag{3.5}$$

where c is to be interpreted as the mean consumption for a given value of y, it being understood that the consumption of individual units deviates from this value by chance.[8] The least squares estimates of α and β (call these a and b), computed from the regression of c on y, are

$$b = \frac{\sum(c - \bar{c})(y - \bar{y})}{\sum(y - \bar{y})^2}, \tag{3.6}$$

$$a = \bar{c} - b\bar{y}, \tag{3.7}$$

where $\bar{c}$ and $\bar{y}$ stand for the mean consumption and income respectively of the group of consumer units, and the summation is over the group. In the numerator of the expression for b, replace y and c by the right hand sides of (3.1) and (3.2), and $\bar{y}$ and $\bar{c}$ by the corresponding sums of means. This gives

$$\begin{aligned} \sum(c - \bar{c})(y - \bar{y}) &= \sum(c_p + c_t - \bar{c}_p - \bar{c}_t)(y_p + y_t - \bar{y}_p - \bar{y}_t) \\ &= \sum(c_p - \bar{c}_p)(y_p - \bar{y}_p) + \sum(c_p - \bar{c}_p)(y_t - \bar{y}_t) \\ &\quad + \sum(c_t - \bar{c}_t)(y_p - \bar{y}_p) + \sum(c_t - \bar{c}_t)(y_t - \bar{y}_t) . \end{aligned} \tag{3.8}$$

From (2.6),

$$c_p = ky_p . \tag{2.6}$$

Inserting (2.6) in (3.8) yields

$$\begin{aligned} \sum(c - \bar{c})(y - \bar{y}) &= k\sum(y_p - \bar{y}_p)^2 + k\sum(y_p - \bar{y}_p)(y_t - \bar{y}_t) \\ &\quad + \frac{1}{k}\sum(c_t - \bar{c}_t)(c_p - \bar{c}_p) + \sum(c_t - \bar{c}_t)(y_t - \bar{y}_t) . \end{aligned} \tag{3.9}$$

[8] On our hypothesis, the relation between the mean value of c and y will be linear only under special conditions. For example, it will be if y_p, y_t, and c_t are distributed according to a trivariate normal distribution. See D. V. Lindley, "Regression Lines and the Linear Functional Relationship," *Journal of the Royal Statistical Society, Supplement*, IX (1947), pp. 218–244.

Given the zero correlations specified in (3.3), the final three terms will differ from zero only because of sampling fluctuations: they will approach zero as the sample size is increased, or average zero over many similar samples. Since our present concern is not with the problem of statistical estimation but with the interpretation of the results, let us suppose the sample to be sufficiently large so that sampling error can be neglected. In that case

$$(3.10) \qquad b = k\frac{\sum(y_p - \bar{y}_p)^2}{\sum(y - \bar{y})^2} = k \cdot P_y,$$

where P_y is the fraction of the total variance of income in the group contributed by the permanent component of income. More generally, of course, b can be regarded as an estimate of the righthand side of (3.10).[9]

The algebraic relation in (3.10) lends itself directly to meaningful interpretation in terms of the permanent income hypothesis. The regression coefficient b measures the difference in consumption associated, on the average, with a one dollar difference between consumer units in measured income. On our hypothesis, the size of this difference in consumption depends on two things; first, how much of the difference in measured income is also a difference in permanent income, since only differences in permanent income are regarded as affecting consumption systematically; second, how much of permanent income is devoted to consumption. P_y measures the first; k, the second; so their product equals b. If P_y is unity, transient factors are either entirely absent or affect the incomes of all members of the group by the same amount; a one dollar difference in measured income means a one dollar difference in permanent income and so produces a difference of k in consumption; b is therefore equal to k. If P_y is zero, there are no differences in permanent income; a one dollar difference in measured income means a one dollar difference in the transitory component of income, which is taken to be uncorrelated with consumption; in consequence, this difference in measured income is associated with no systematic difference in consumption; b is therefore zero. As this explanation suggests, P_y, though *defined* by the ratio of the variance of the permanent component of income to the variance of total income, can be *interpreted* as the fraction of any difference in measured income that on the average is contributed by a difference in the permanent component. This point is developed more fully below.

Substitute (3.10) in (3.7), replace $\bar{c}$ by $\bar{c}_p + \bar{c}_t$, $\bar{y}$ by $\bar{y}_p + \bar{y}_t$, and

[9] In the special case of the preceding footnote, $\beta = kP_y$.

$\bar{c}_p$ by $k\bar{y}_p$. The resulting expression can then be written:

(3.11) $$a = \bar{c}_t - kP_y\bar{y}_t + k(1 - P_y)\bar{y}_p\,.$$

The elasticity of consumption with respect to income at the point (c, y) is

(3.12) $$\eta_{cy} = \frac{dc}{dy}\cdot\frac{y}{c} = b\cdot\frac{y}{c} = kP_y\cdot\frac{y}{c}\,.$$

Suppose that the mean transitory components of both income and consumption are equal to zero, so that $\bar{y} = \bar{y}_p$, $\bar{c} = \bar{c}_p$. In this special case

(3.13) $$\frac{\bar{y}}{\bar{c}} = \frac{1}{k}\,.$$

It follows that if the elasticity is computed at the point corresponding to the sample mean:

(3.14) $$\eta_{cy} = P_y\,.$$

Consider, now, the regression of y on c, say

(3.15) $$y = a' + b'c\,.$$

By the same reasoning it can be shown that, sampling errors aside,

(3.16) $$b' = \frac{1}{k}P_c\,,$$

where P_c is the fraction of the variance of consumption contributed by the permanent component, and

(3.17) $$a' = \bar{y}_t - \frac{1}{k}P_c\bar{c}_t + \frac{1}{k}(1 - P_c)\bar{c}_p\,.$$

The elasticity of consumption with respect to income computed from this regression is

(3.18) $$\eta'_{cy} = \frac{dc}{dy}\cdot\frac{y}{c} = \frac{1}{b'}\cdot\frac{y}{c} = \frac{k}{P_c}\cdot\frac{y}{c}\,.$$

Again, if $\bar{y}_t = \bar{c}_t = 0$,

(3.19) $$\eta'_{cy} = \frac{1}{P_c}\,,$$

if evaluated at the point corresponding to the sample mean.[10]

[10] For the logarithmic alternative described by (2.6′), (3.1′), (3.2′), and (3.3′) the analogues to the results given in the text are

(3.10′) $$B = P_Y,$$

(3.11′) $$A = K + \bar{C}_t - \bar{Y}_tP_Y + \bar{Y}_P(1 - P_Y),$$

Some of these results are presented in graphic form in Figure 3 for the special case in which the mean transitory components of income and consumption are zero.

Consider the consumer units with a particular measured income, say y_0, which is above the mean measured income for the group as a whole. Given zero correlation between the permanent and transitory components of income, the average permanent income of these units

FIGURE 3

Hypothetical Relation between Measured Consumption and Measured Income (mean transitory components equal zero)

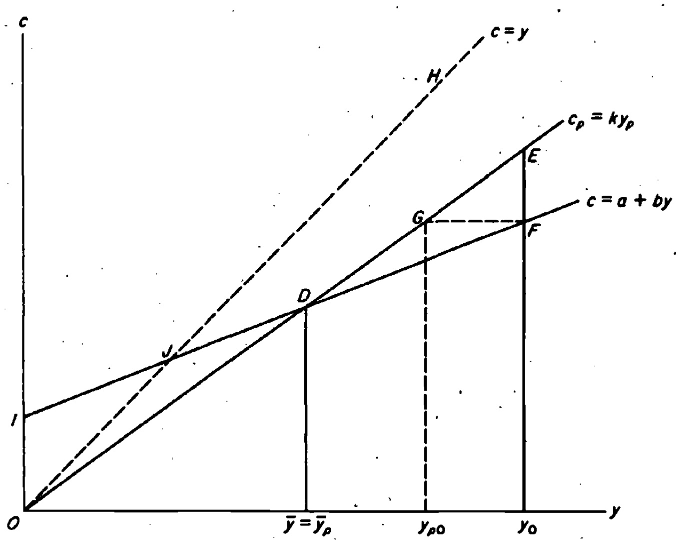

is less than y_0; that is, the average transitory component is positive. These units have been classified together precisely because their measured income is a certain amount *above* the average income of the group; such a relatively high measured income could have been

---

(3.12′) $$\eta_{cy} = \frac{dC}{dY} = B = P_Y,$$

(3.16′) $$B' = P_C,$$

(3.17′) $$A' = -K + \bar{Y}_t - \bar{C}_t P_C + \bar{C}_p(1 - P_C),$$

(3.19′) $$\eta'_{cy} = \frac{1}{P_C}.$$

These results are in some ways simpler and more appealing than those in the text, since the elasticity of consumption with respect to income is the same everywhere and hence equal to P_Y or $1/P_C$ without the necessity of assuming the mean transitory components to be zero.

received despite unfavorable transitory effects; clearly, it is more likely to have been received because of favorable transitory effects; the winners in any particular set of races may well be better on the average than the losers but they are also likely to have had more than their share of good luck. Put more rigorously, the transitory component of income is positively correlated with the sum of itself and a variable (the permanent component) with which it is itself uncorrelated.[11] What about the average transitory component of consumption for these units? The corresponding component of their income is positive because the transitory component of income helped to determine which units had a measured income of y_0 and so were classified together; given, however, that the transitory components of income and consumption are uncorrelated, a classification by income is random with respect to transitory components of consumption; in consequence, the latter tend to average out to the average for the group as a whole, which is assumed to be zero. The average consumption of units with a measured income y_0 is therefore equal to their average permanent consumption. On our hypothesis, this is k times their average permanent income. If y_0 were not only the measured income of these units but also their permanent income, their mean consumption would be ky_0 or y_0E. Since their mean permanent income is less than their measured income, their average consumption, y_0F, is less than y_0E.

By the same reasoning, for consumer units with an income equal to the mean of the group as a whole, or $\bar{y}$, the average transitory component of income as well as of consumption is zero, so the ordinate of the regression line is equal to the ordinate of the line OE which gives the relation between permanent consumption and permanent income. For units with an income below the mean, the average transitory component of income is negative, so average measured consumption is greater than the ordinate of OE. The regression line therefore intersects OE at D, is above it to the left of D, and below it to the right of D.

Let us return to the income class y_0. Draw a horizontal line through F. The abscissa of the point G, where this line intersects OE, is the permanent income associated with a permanent consumption y_0F. This income, labelled y_{p0} on the figure, is therefore the average permanent component of the income of the members of the income class y_0, and $(y_0 - y_{p0})/(y_0 - \bar{y})$ is the fraction of the deviation of their average income from the average for the group attributable to the transitory component. If this fraction is the same for all income classes, IF is a straight line, and the common value of the fraction is

[11] See Friedman and Kuznets, *op. cit.*, pp. 327–332, esp. footnotes 10 and 13.

$1 - P_y$.[12] The higher this fraction, the flatter *IF* and conversely. At one extreme, if P_y is zero, that is, if all members of the group have the same permanent component, average consumption is the same for all income classes and *IF* is horizontal. At the other extreme, if $P_y = 1$, so transitory components are all zero, *IF* coincides with *OE*.

If k is less than unity, permanent consumption is always less than permanent income. As is clear from the figure, however, it does not follow that measured consumption is necessarily less than measured income. The line *OH* on the figure is a 45 degree line along which $c = y$. The vertical distance between this line and *IF* is average measured savings. Point *J* is the "break even" point at which average measured savings are zero. To the left of *J*, average measured savings are negative, to the right, positive; as measured income increases, so does the ratio of average measured savings to measured income. Our hypothesis thus yields a relation between measured consumption and measured income that reproduces the broadest features of the corresponding regressions that have been computed from observed data.

For the special case for which Figure 3 is drawn, k could be readily computed from observed data on the measured consumption and measured income of a group of consumer units, since average measured consumption and average measured income then equal the corresponding average permanent components. The line *OE* in the figure therefore goes through the point describing the mean income and consumption of the group, so $k = \bar{c}/\bar{y}$. P_y could then be computed from the relation between the regression of c on y (the line *IF*) and the line *OE*, and P_c from the corresponding relation between the regression of y on c and the line *OE*.[13]

If the mean transitory component of consumption is not zero, the curve *IF* is shifted vertically by a corresponding amount—upwards, if the mean transitory component is positive, downwards, if it is negative. Clearly, there is no way of distinguishing such a shift from a change in k. Similarly, a positive mean transitory component of income shifts *IF* to the right, a negative mean, to the left. For a straight line, there is no way of distinguishing such horizontal shifts from vertical shifts produced by a mean transitory component of

[12] See *ibid.* pp. 332–336, 358. Figure 3 is essentially the same as Chart 28 on p. 333.

[13] The estimation problem is the classical one of "mutual regression" or regression when "both variables are subject to error." See D. V. Lindley, *op. cit.*, for an excellent analysis of the problem and survey of the literature. Many of our equations duplicate equations in his paper. As Lindley points out, there are no efficient statistics for estimating all the parameters in the model from sample data. The method described in the text is therefore not statistically efficient. The usual solution is to assume the ratio of the variance of y_t to the variance of c known, in which case efficient statistical procedures do exist.

consumption. It follows that, if the mean transitory components cannot be set equal to zero, data for one group for one time period are inadequate to estimate all the parameters. Some other source of information is required as well.

Our hypothesis gives a major role to certain features of the income distribution generally neglected in consumption studies. It asserts that some of the most strikingly uniform characteristics of computed regressions between consumption and income are simply a reflection of the inadequacy of measured income as an indicator of long-run income status. In consequence, differences among various groups of consumer units in observed marginal propensities to consume may not reflect differences in underlying preferences for consumption and wealth at all; they may reflect primarily the different strength of random forces, including errors of measurement, in determining measured income. Fortunately, considerable evidence is available on the importance of transitory components of income from studies of changes over time in the relative income status of individuals or consumer units. One of the attractive features of our hypothesis is that it enables us to bring this independent body of evidence to bear on the interpretation of consumption behavior; such evidence can provide some of the additional information required when transitory components of income and consumption cannot be supposed to be zero.

Before examining these data, however, we shall first examine the consistency of the hypothesis with some of the major general findings of empirical studies of consumption behavior and its relation to the relative income hypothesis suggested by Brady and Friedman, Duesenberry, and Modigliani. This will serve the double purpose of bringing out more fully the implications of the hypothesis and of suggesting the evidence that recommends its acceptance as a provisional working hypothesis.